M&SOM MANUFACTURING & SERVICE OPERATIONS MANAGEMENT
Vol. 11, No. 1, Winter 2009, pp. 109–127
ISSN 1523-4614 | EISSN 1526-5498 | 09 | 1101 | 0109

informs®

DOI 10.1287/msom.1070.0199

Long-Term Contracts Under the Threat of Supplier Default

Robert Swinney, Serguei Netessine
The Wharton School, University of Pennsylvania, Philadelphia, Pennsylvania 19104
{rswinney@wharton.upenn.edu, netessin@wharton.upenn.edu}

Contracting with suppliers prone to default is an increasingly common problem in some industries, particularly automotive manufacturing. We model this phenomenon as a two-period contracting game with two identical suppliers, a single buyer, deterministic demand, and uncertain production costs. The suppliers are distressed at the start of the game and do not have access to external sources of capital; hence, revenues from the buyer are crucial in determining whether default occurs. The production cost of each supplier is the sum of two stochastic components: a common term that is identical for both suppliers (representing raw materials costs, design specifications, etc.) and an idiosyncratic term that is unique to a given supplier (representing inherent firm capability). The buyer chooses a supplier and then decides on a single- or two-period contract. Comparing models with and without the possibility of default, we find that, without the possibility of supplier failure, the buyer always prefers short-term contracts over long-term contracts, whereas this preference is typically reversed in the presence of failure. Neither of these contracts coordinates the supply chain. We also consider dynamic contracts, in which the contract price is partially tied to some index representing the common component of production costs (e.g., commodity prices of raw materials such as steel or oil), allowing the buyer to shoulder some of the risk from cost uncertainty. We find that dynamic long-term contracts allow the buyer to coordinate the supply chain in the presence of default risk. We also demonstrate that our results continue to hold under a variety of alternative assumptions, including stochastic demand, allowing the buyer the option of subsidizing a bankrupt supplier via a contingent transfer payment or loan and allowing the buyer to unilaterally renegotiate contracts. We conclude that the possibility of supplier default offers a new reason to prefer long-term contracts over short-term contracts.

Key words: supply chain management; financial default; contracting; game theory
History: Received: September 2, 2006; accepted: August 13, 2007. Published online in *Articles in Advance* January 25, 2008.

1. Introduction

The relationship between manufacturers and suppliers in the American automotive industry is not always a cooperative one. American carmakers are on a perpetual quest to match the procurement costs of their competitors by increasing supply chain efficiency. Throughout the 1990s, an increased awareness of the value of cooperative buyer-supplier relationships sparked interest in fostering strategic partnerships between automotive manufacturers and suppliers. Nevertheless, despite the perceived value of collaborative behavior with suppliers, auto manufacturers often engage in adversarial and caustic supplier management tactics, typically employing one tool more than any other: direct pressure to reduce procurement prices.

In the 1990s, for example, Ford famously dictated an across-the-board 5% price decrease to all of its suppliers (Stallkamp 2005). In 2005, Lear, a key seat supplier to Chrysler, attempted to negotiate higher prices to cover recent sharp cost increases that plagued the automotive supplier base. When Lear threatened to cease shipping products to Chrysler, the automaker promptly took the supplier to court to enforce the terms of their contract, despite the fact that Lear posted a net loss of nearly $600 million in the fourth quarter of 2005 alone. This emphasis on low procurement prices has had a clear adverse effect on the suppliers. Profit margins are low across the industry, and many suppliers routinely lose money (Wernle 2006a, Wynn 2006). In 2005, Delphi, the largest supplier of automotive parts in the country, was in bankruptcy, as were numerous smaller suppliers.

While the poor financial health of the suppliers is exacerbated by the low procurement costs demanded by automakers in response to the competitive North American automotive landscape, several additional factors have played a key role in the current perilous state of the automotive supplier base. Bo Andersson, General Motors Vice President for Global Purchasing and Supply Chain, cites four critical issues: production cost increases, unstable domestic volume, legacy pension plans (resulting in larger overhead expenses), and difficult access to capital (Andersson 2006). Our analysis will touch on each of these key factors.

From a buyer's point of view, losing a supplier to bankruptcy can have various consequences: At one extreme, if the supplier ceases operations, the buyer may have to switch to a new supplier (possibly at a higher cost), whereas, at another extreme, if the supplier continues normal operations without disruption, the buyer may have to help support the supplier financially, as was the case in some of the most visible supplier bankruptcies of 2005 (Delphi; see Nussel and Barkholz 2006). When interior parts supplier Collins & Aikman declared bankruptcy in 2005, automakers had little choice but to sustain the supplier—roughly 90% of vehicles made in North America have at least one component produced by the supplier—and the estimated total cost to the Big Three auto manufacturers was $532 million, resulting from the cancellation of loan repayments, parts price increases, operating subsidies, and legal fees (Barkholz and Sherefkin 2007).

Thus, the possibility of losing a supplier to failure may affect the decisions that a buyer makes, including the price and the length of contracts. Should a buyer pay more to avoid losing a supplier to bankruptcy, or should the buyer pay less because the supplier is risky? Should the buyer make a long-term commitment to the supplier, or should the firm minimize the length of its exposure to a risky partner? Should buyers bear some of the risk of cost uncertainty by compensating suppliers in a dynamic manner? These are the trade-offs we seek to capture, via an operationalized model of buyer-supplier relations under the threat of supplier failure. The motivating example is the automotive industry, but the model is relevant to many scenarios: supply chains with members in financial distress, supply chains with start-ups prone to bankruptcy, etc. Based on these examples from the automotive industry, we believe that these types of relationships have several key characteristics.

1.1. Uncertain Supplier Production Costs

Tooling and capacity installation leadtimes tend to be long. Hence, the buyer must commit to the supplier well in advance of the finalized design of the component. In addition, raw materials costs are uncertain and often have a large impact on the supplier's margins; Standard & Poor's June 2006 industry survey of autos and auto parts cites high raw materials costs as a key factor in the current distressed state of the supplier base (Levy and Ferazani 2006). Thus, the actual per-unit production cost to the supplier is unknown at the time of contracting. We assume that this risk is not diversifiable because it derives from volatility not captured in any current futures market (e.g., outputs from higher-tier suppliers or inherent uncertainty in design and production techniques). Even risk in raw materials prices cannot always be hedged—Sherefkin (2006) describes how Ford has struggled to create a futures market for automotive sheet steel.

1.2. Strong Bargaining Power of the Buyer

The buyer tends to be a large firm with most of the bargaining power, whereas the supplier is the smaller firm at risk of default. For example, in the American auto industry, parts suppliers have few potential buyers and face weak bargaining positions and low profit margins, whereas in Europe the situation is far less bleak for the suppliers (Lewin 2006). Hence, the buyer has most of the bargaining power and offers the contract to the supplier, choosing both price and the length of the contract.

1.3. Extended Sales Horizons

The product being supplied will be used over several sales periods (e.g., a particular model of a car tends to be sold for five to seven years before a major redesign). Because the cost of switching suppliers is typically high (due to large asset-specific investments made by the manufacturers into a particular supplier), the financial health of a dedicated supplier over the sales horizon is critical to the buyer. Furthermore, the buyer has the option of contracting with a supplier for a single sale period or for multiple periods; both practices are observed in the auto industry (Dyer 1996).

These three points form the core of our model. Here we will analyze and evaluate the performance of both long- and short-term contracts when suppliers face a risk of failure and determine under what conditions a particular contract type is preferred. We compare these results to a model with no supplier failure in which short-term contracts are always preferred and find that under the threat of supplier default this preference is typically reversed. We further consider dynamic contracts, which partially compensate the suppliers for the realized value of production costs, and find that long-term dynamic contracts perform better than static contracts and are capable of achieving the centralized system optimal profit.

Our results complement existing literature on the value of long-term contracts in supply chains by demonstrating their advantages despite controlling for many of the typical reasons a buyer would have to prefer these contracts. For example, management literature often discusses long-term relationships or contracts as a method of developing trust between buyers and suppliers, but we do not consider trust issues. Indeed, long-term contracts are preferred even though the buyer is completely self-interested. Similarly, long-term relationships are known to induce otherwise unsupportable actions in repeated games, thus increasing their value relative to short-term relationships. In our model, the value comes not from inducing actions but rather from reducing the expected cost of supplier default to the buyer. Overall, we offer a new reason to use long-term contracts: to reduce the damage from supplier default. Our finding is consistent with current practices in the Japanese auto industry, in which long-term relationships are common and supplier defaults are, relative to the Untied States, uncommon.

The remainder of this paper is organized as follows: The next section provides a brief literature review, and §3 describes the model. Section 4 analyzes a benchmark model with suppliers that never enter bankruptcy, and §5 considers suppliers prone to bankruptcy. Section 6 explores a class of contracts that partially compensate the suppliers for production costs in a dynamic fashion. Section 7 discusses three interesting extensions (demand uncertainty, contingent transfer payments, and normally distributed costs), and §8 concludes the paper.

2. Literature Review

Three broad areas of research are relevant to our paper. The first focuses on understanding the nature of buyer-supplier relationships in the auto industry. The second is primarily theoretical, concerning such issues as repeated contracting between firms, relational agreements, and contracting under cost uncertainty. The third area of research focuses on the effects of financial distress and supplier disruption risk.

There is a strong tradition, particularly in management literature, of research into the types of relationships that exist between buyers and suppliers in the automotive industry. Tang (1999) provides an excellent discussion. Much of this work focuses on the historical and current differences between Japanese-style and American-style supply networks; see, for example, McMillan (1990), Dyer et al. (1998), Dyer (1996), and references therein. The traditional Japanese networks of suppliers, called *keiretsu*, are markedly different from their American counterparts. The former are characterized by fewer suppliers, long-term relationships, and heavy cooperation, whereas the latter traditionally used a large number of small suppliers, short-term contracts, and non cooperative or adversarial behavior. Japanese firms are often heavily invested in their suppliers, wholly or partly owning their closest partners in many instances (Dyer et al. 1998). Even when buyers are not directly invested in suppliers, there is indirect investment via the value of the ongoing relationship. Such close networks of suppliers inextricably link the financial health of firms in the supply chain.

Empirically, there is evidence both that the financial health of suppliers matters to buyers and that collaborative relationships are beneficial to suppliers. Choi and Hartley (1996), in a survey of suppliers and automakers in the U.S. industry, find that financial issues are a primary factor in supplier selection and that greater importance is placed on financial health by downstream firms (i.e., the auto assemblers) than by upstream firms. Srinivasan and Brush (2006) find empirical evidence (outside of the auto industry) that buyer-supplier collaboration and target pricing benefit the financial performance of suppliers.

In recent years American companies have attempted to emulate some aspects of the Japanese system, in particular the narrowing of supplier bases and longer-term contracts; see, for example, Dyer (1996), Tang (1999), and more recently Wernle (2006b). Nevertheless, as Sako and Helper (1998) demonstrate, trust between supply chain partners is still higher among Japanese firms than among American firms. Furthermore, Rudambi and Helper (1998) find empirical evidence for noncooperative behavior in the U.S. auto industry, suggesting that significant differences still persist between the American and Japanese auto industries. Consequently, a buyer has incentive to keep suppliers in good financial standing only if it benefits the buyer through lower costs. This is the situation that we model: a noncooperative supplier–buyer relationship in which the buyer is concerned with the failure of a supplier only to the extent that it might be costly for the buyer to switch suppliers.

We examine two scenarios: a long-term contract that covers the entire horizon and a series of repeated, short-term contracts. Li and Debo (2007) examine a similar model, quantifying the value of commitment to a single firm in a two-period newsvendor context in which suppliers have private information about their production costs. The buyer runs an auction to pick a supplier and may choose to run an auction in each period or to commit to a supplier in the first period. They find that a long-term contract increases the aggressiveness of supplier bidding and thus helps to counteract the effect of information asymmetry. Whereas the results of Li and Debo (2007) are driven by information asymmetries, our model is driven by failure risk and cost uncertainty; ex ante, there is no information asymmetry in our setting.

The short-term contracts in our model are essentially relational, which constitute an emerging topic in operations literature. The term relational contract refers to the fact that the enforcement of the contract comes from the value of the future relationship rather than from direct legal enforcement. Examples of related papers include Taylor and Plambeck (2007) and Atkins et al. (2005). These models contain single-buyer, single-supplier relationships, in contrast to our model with two suppliers. Tunca and Zenios (2006), on the other hand, examine a model of repeated contracting with multiple suppliers and buyers. The authors compare relational contracts for high-quality components with and without a secondary electronic market for low-quality components. Their model, in contrast to ours, has a powerful high-quality supplier who offers the contract to the buyer and does so before a group of low-quality suppliers (i.e., the high-quality supplier is a Stackelberg leader). None of these papers considers the endogenous effect of one partner leaving the relationship because of failure, high production costs, low capital, etc. (although discount factors in repeated games can be thought of as an exogenous probability of relationship termination).

There is also related literature comparing long- and short-term contracts under cost uncertainty, surveyed by Kleindorfer and Wu (2003). In much of this literature, long-term contracts refer to those signed prior to the realization of some random variable (e.g., cost or demand) and short-term contracts are those signed after the realization of this stochastic event, with demand usually occurring in a single period. One exception is the multiperiod model of Cohen and Agrawal (1999); however, in contrast to our model, contract price is not a decision variable (although contract length is), and there is no chance of contract termination due to supplier default. For a general reference on contracting in supply chains, we refer to the survey by Cachon (2003).

Papers that have considered firm bankruptcy or failure in an operational context are Archibald et al. (2002) and Swinney et al. (2005). Both of these papers define failure in the same way as the present work: If capital falls below a prespecified level, the firm ceases to exist. The first paper looks at a monopoly setting, and the latter considers duopolies. Neither considers contracting effects between firms.

In a closely related paper, Babich et al. (2007) analyze a model with multiple suppliers and a single buyer, wherein suppliers face an exogenous probability of default and act as Stackelberg leaders in setting the wholesale price for a downstream newsvendor. Our model differs in that the default risk is endogenous (i.e., it is a function of the contract price between buyer and supplier, implying that the buyer's business has a significant effect on the supplier's financial health) and the bargaining power lies with the buyer.

Furthermore, Babich et al. (2007) consider a single period model, whereas our model considers contracting effects over multiple periods. Related work includes Babich (2006).

Tomlin (2006) considers methods of mitigating disruption risk when the buyer contracts with reliable or unreliable suppliers. The focus is on sourcing and contingency strategies to help mitigate the effects of disruption risk under uncertain demand. We focus on contract parameters that directly minimize the buyer's expenses from supplier failure under uncertain costs. In much of the disruption risk literature the risk of failure or default is exogenous, whereas in our analysis it is endogenously determined by the contract price between the buyer and the supplier. To summarize, we are not aware of any work that, like ours, uses contracting to mitigate the harmful effects of endogenous supplier default.

3. The Model

There are two identical suppliers (subscript $i = 1, 2$) and a single buyer (subscript b). The analysis is unchanged if we consider a pool of an arbitrary number of potential (ex ante identical) suppliers. The buyer requires some critical component in each of two periods and will contract with one supplier at a time to obtain the component. We assume that dual sourcing is too costly to be considered (i.e., there is a large fixed cost to doing business with any supplier); Tang (1999) describes how American companies have recently reduced supplier bases and increased the frequency of sole-sourcing to control costs. Demand is identical and known in each period and without loss of generality is normalized to one, consistent with the automotive industry, in which short-term forecasts of sales (and, in particular, of procurement quantities from suppliers) are fairly accurate for mature products. (We will relax the assumption of deterministic demand in §7.1.) The buyer sells the finished product for price s.

Each supplier has a linear unit production cost that is the sum of two independent stochastic components: common costs c_t, $t = 1$, 2, which are identical for both suppliers but may be time-dependent, and idiosyncratic costs d_i, $i = 1$, 2, which are unique to a given supplier but may be correlated with one another and are constant across time. Thus, the total production cost of supplier i at time t is $c_t + d_i$. For generality, we allow both random variables to have support in $(-\infty, \infty)$, though they may be restricted to some smaller interval. The uncertainty in common cost arises from stochastic elements that affect both firms for example, raw materials costs or product design specifications. The uncertainty in a firm's idiosyncratic cost arises from factors unique to a given supplier, such as the efficiency of a supplier's production facilities or its level of technical expertise, reflecting an implicit notion of cost discovery in the production process; suppliers may have some estimate about their inherent efficiency, but, until they physically produce a large number of units, the precise value of this cost is unknown.

All firms have identical beliefs that c_1 has distribution $F(\cdot)$, unimodal density $f(\cdot)$, and finite mean $\mu_1 > 0$. The second period common cost may depend on the realized value of the first period cost. The conditional mean is defined as $\mu_2(x) = \mathbb{E}(c_2 \mid c_1 = x)$, the unconditional mean is $\mu_2 > 0$, and both are assumed to be finite.

The idiosyncratic costs of the two suppliers may be correlated; however, because the suppliers are assumed to be ex ante identical, the marginal distributions are identical as well, and thus all firms believe that the marginal cdf is $G(x)$ and the marginal pdf is $g(x)$. The conditional mean is $\mu_d(x) = \mathbb{E}(d_2 \mid d_1 = x)$, and the unconditional mean is $\mu_d > 0$, where both are assumed to be finite. The correlation coefficient between d_1 and d_2 is denoted ρ_d. For technical reasons, we assume that $x - \mu_d(x)$ is monotonically increasing. (The complementary case of $x - \mu_d(x)$ monotonically decreasing in x is impossible if the marginal distributions of d_1 and d_2 are identical.) For example, this assumption holds if d_1 and d_2 are bivariate normal.

There is no private information in the model; all parties learn the values of all random variables when they are realized. For example, if supplier 1 produces in the first period, then by the start of the second period both suppliers and the buyer know the values of c_1 and d_1. The second period common cost and d_2 are still unknown at this point, though the realized value of d_1 may convey some information about d_2 if the two have non-zero correlation. We assume no private information for a variety of reasons reflective of the automotive industry: For example, sources of common cost uncertainty are clearly known well by

Table 1 **Cash Flows to the Suppliers in Each Period**

In	Out
Existing capital	Interest payment
Loans	Fixed operating expenses
Revenue from buyer	Production expenses

the buyer (e.g., raw materials, design specifications), while the buyer is also likely to have a good estimate of the idiosyncratic capabilities of each supplier from existing (or previous) relationships. (Given the limited number of suppliers in the auto industry, it is unlikely that a buyer has never worked with a given supplier in the past). Finally, most suppliers are public companies, and much of the labor is unionized, meaning that the general financial state of each supplier and labor costs are public information.

We assume that both suppliers are distressed at the start of the game; that is, they are already in danger of bankruptcy when the buyer offers a contract at the start of period one. Suppliers are already heavily leveraged at the time of contract signing and hence cannot borrow additional funds from external sources. (It may be desirable to borrow or transfer funds from the buyer; this is discussed in §7.2). Payments on outstanding debt are made in alignment with the production time periods of our model. Cash flows to the suppliers are depicted in Table 1. The total capital level of a supplier is the difference between the cash inflows and the cash outflows. If the total cash flow becomes negative, then the supplier is incapable of making the necessary debt payments and will enter bankruptcy. (Throughout, we use the terms failure, default, and bankruptcy synonymously for the sake of variety.) Without loss of generality, we normalize the sum of the first two components in each column to zero; that is, existing capital plus loans less interest payments and fixed operating expenses equal zero. Consequently, the supplier fails if, at the end of any period, the total profit from operations with the buyer falls below the bankruptcy threshold of zero.

If the supplier fails, then the relationship with the buyer is broken, and the buyer must switch to a new supplier incurring a switching cost of $k \geq 0$. If the buyer chooses to switch suppliers after one period of production, the same switching cost k is incurred. Alternatively, k may be interpreted as a fixed setup cost incurred upon doing business with any supplier. This value also includes any additional costs incurred by the buyer to expedite production with the alternative supplier.

The expected profit to the buyer is π_b, and the expected profit to supplier i is π_i. We assume that the buyer maximizes expected profit and that supplier i agrees to any contract with expected profit no less than the reservation level of zero. Non-zero reservation levels and bankruptcy thresholds merely add constant terms to the contract parameters and do not alter the qualitative nature of the results, and so for notational simplicity they are omitted.

We compare two types of contracts: short- and long-term. The short-term contract (Figure 1, top) provides

Figure 1 **Sequence of Events in Short-Term (Top) and Long-Term (Bottom) Contracts**

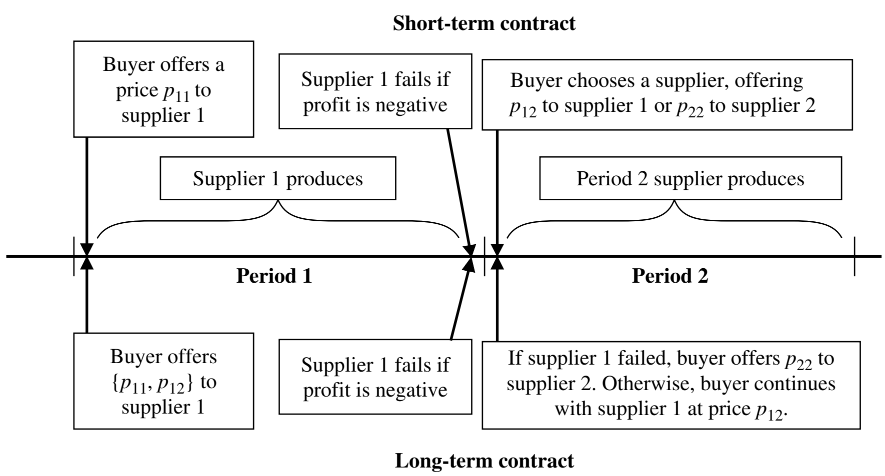

legally enforceable terms for only one period of procurement, whereas the long-term contract (Figure 1, bottom) covers both periods. Because suppliers are ex ante identical, without loss of generality, we assume that the buyer contracts first with supplier 1. The price offered to supplier i in period t is denoted p_{it}; thus, the relevant prices are p_{11} and p_{12} (respectively, first- and second-period prices with supplier 1) and p_{22} (second period price with supplier 2, if supplier 1 fails or if the buyer chooses to switch suppliers). We denote the optimal profits and prices in a long-term contract with the superscript l and the optimal prices and profits in a short-term contract with the superscript s. In all contracts we seek prices that are subgame perfect—that is, the offered price at any given time must be optimal from the point of view of the buyer at that point in time. Initially, we assume that all contracts are static (i.e., the prices cannot depend on the realized value of any random variables), because this is commonly the situation observed in the U.S. auto industry (McMillan 1990). Less rigid contract forms are analyzed in §§6 and 7.

One might ask why the long-term contract is not subject to renegotiation at the start of the second period, provided the original supplier does not fail. Indeed, the susceptibility of long-term contracts to renegotiation may have a large effect on the performance of the contract and supply chain coordination; see, for example, Plambeck and Taylor (2007). There are two relevant types of renegotiation to consider: that initiated by the supplier (i.e., hold-up) and that initiated by buyer (exploiting their powerful bargaining position). We discuss buyer-initiated renegotiation in §7.3. As for supplier renegotiation, we explicitly exclude this possibility on the assumption that the buyer is powerful enough to thwart any attempt by a supplier at hold-up. A relevant and timely example is that of the Lear Corporation (Barkholz and Sherefkin 2006), in which Chrysler took Lear to court to enforce the contractual terms despite ample evidence of Lear's financial distress. In response to Lear's attempts to raise prices, Chrysler replaced the firm with rival Johnson Controls as seat supplier for the Dodge Ram starting in 2008, a move that further imperiled Lear's financial health. Thus, in the case of Lear, there are both court-enforceable and relational effects in play: The court enforced the formal contract with Chrysler, and Chrysler initiated a relational punishment by switching suppliers on a later car model.

4. A Benchmark Model Without Failure

We first examine a model in which suppliers never default (i.e., negative profits do not force the buyer to switch suppliers). By comparing the results to a model with supplier failure, we will determine how the threat of default alters the behavior of the buyer and the suppliers. This model also serves as a benchmark, providing the maximum expected profit that the coordinated system can achieve.

The primary difference between the model with failure and the model without failure is that, in the latter, the buyer never switches suppliers in a two-period contract (and hence never incurs a switching cost). In a single-period contract, the buyer switches suppliers only if it is cost-effective to do so after taking into account the switching costs. Consequently, the model without failure provides an upper bound on the performance of the system with failure. This is an important observation, because we will derive a contract in a later section that achieves this upper bound even in the presence of suppliers prone to default.

The following theorem details the optimal contracts (both short- and long-term) for the no-failure case. It is useful to define the following critical cost value: $\alpha = \{x\colon k + \mu_d(x) = x\}$, if such a solution exists, otherwise $\alpha = \infty$. Given that $x - \mu_d(x)$ is assumed to be increasing in x, there is at most one solution to this equation.

Theorem 1. (i) *In the absence of failure, the optimal short-term contract is* $p^s_{11} = \mu_d + \mu_1$, $p^s_{12} = d_1 + \mu_2(c_1)$, *and* $p^s_{22} = \mu_d(d_1) + \mu_2(c_1)$, *where* d_1 *is the realized value of supplier 1's idiosyncratic costs. The buyer switches suppliers in the second period if* $k + \mu_d(d_1) < d_1$ *(i.e., if* $\alpha < d_1$*). The resulting expected profit for the buyer is*

$$\pi^s_b = 2s - \mu_d - \mu_1 - \mu_2 - \mathbb{E}\min(k + \mu_d(d_1), d_1). \quad (1)$$

(ii) *In the absence of failure, the optimal long-term contract is any pair* $\{p^l_{11}, p^l_{12}\}$ *such that* $p^l_{11} + p^l_{12} = 2\mu_d + \mu_1 + \mu_2$, *and the resulting expected profit for the buyer is*

$$\pi^l_b = 2s - 2\mu_d - \mu_1 - \mu_2. \quad (2)$$

(iii) *In the absence of failure, the buyer always prefers a short-term contract to a long-term contract. Furthermore, among single-sourcing contracts, the optimal short-term contract achieves the centralized system optimal profit, which we denote by* $\bar{\pi}_b$.

Proof. All proofs appear in the technical appendix (online). □

The reason that the buyer prefers the short-term contract is simple: When there is no possibility of failure, the buyer switches suppliers only when the alternative supplier has lower expected costs. The long-term contract eliminates the buyer's opportunity to switch suppliers, an option that always has value in the absence of supplier default.

The fact that the short-term contract achieves the system optimal (first best) profit is also intuitive. The short-term contract in Theorem 1 ensures that the buyer uses the most efficient supplier (after accounting for switching costs) in each period. This is the same goal of the centralized system (i.e., if one firm controlled the buyer and both suppliers). Thus, in a short-term contract without failure, the total profit in the system is maximized, and the contract is optimal among all (single-sourcing) contract types. Figure 2 details the second-period actions of the buyer as a function of the realized values of the random variables.

5. Suppliers Prone to Default

In this section we consider contracts between a buyer and suppliers that are prone to default. Recall that suppliers default if, at the end of any period, their total profit is negative. For the purposes of the buyer, default matters only if it happens at the end of period 1, i.e., if $p_{11} < c_1 + d_1$; otherwise, supplier 1 survives.

5.1. Short-Term Contract

Under a short-term contract, the buyer may switch suppliers voluntarily (as in the no-failure case) or involuntarily (if supplier 1 defaults). In determining the optimal second-period action, the buyer faces one of three separate cases, depending on the realized values of the cost parameters. In the first case (which we call Region I), supplier 1 survives the first period and the buyer chooses to continue with that supplier in the second period. This case occurs if total first-period costs are small (i.e., below p_{11}) and if supplier 1's idiosyncratic costs are low. In the second case (Region II), supplier 1 survives, but the buyer switches to supplier 2 in the second period. This occurs if total first-period costs are small but supplier 1's idiosyncratic costs are large. In the third case (Region III), supplier 1 defaults, and the buyer must switch to supplier 2 in the second period. This happens if total first-period costs are high (i.e., above p_{11}). The buyer controls the size and shape of these regions via the offered prices; see Figure 3 for an illustration. Comparing this graph to Figure 2, we see that the addition of Region III forces the buyer to switch suppliers for a much larger region of the probability space.

Figure 2 Optimal Second-Period Action of the Buyer as a Function of the Realized Values of c_1 and d_1 in the Short-Term Contract When There Is No Possibility of Supplier Failure

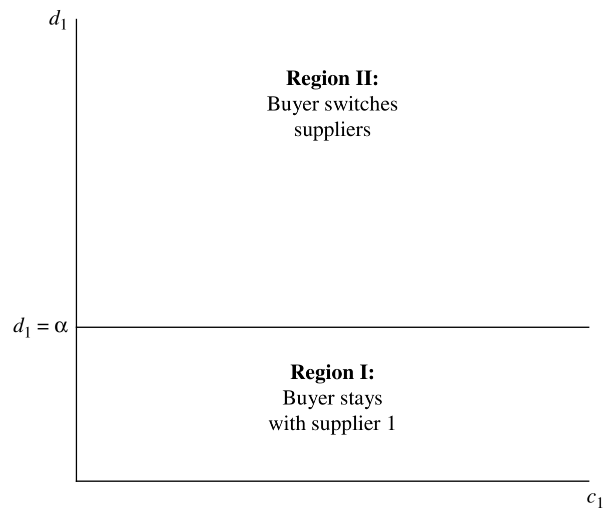

The following lemma details when these regions occur and characterizes the resulting optimal short-term contract. We adopt the convention that $\pi_b^s(p_{11})$ denotes the buyer's optimal expected profit in a short-term contract as a function of p_{11} (e.g., with all other prices set optimally), and hence $\pi_b^s = \max_{p_{11}} \pi_b^s(p_{11})$.

Lemma 1. *Define* p_{11}^* *as the solution to*

$$-1 + \int_{-\infty}^{\alpha} (\mu_d(x) + k - x) f(p_{11}^* - x) g(x)\, dx = 0.$$

Figure 3 Optimal Second-Period Action of the Buyer as a Function of the Realized Values of c_1 and d_1 in the Short-Term Contract When There Is a Risk of Supplier Failure

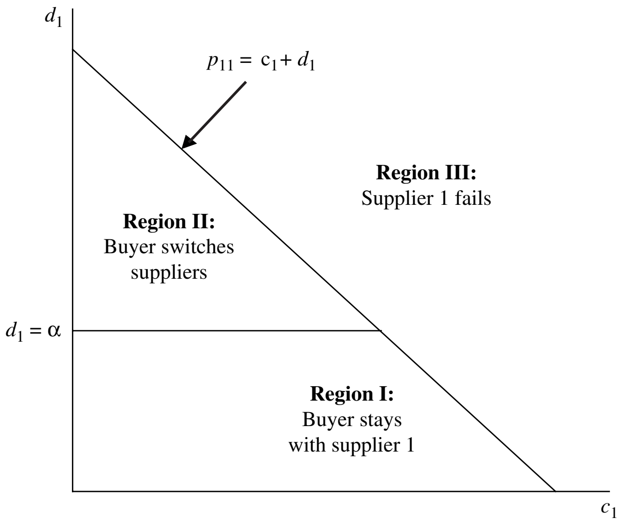

Then the buyer's optimal short-term contract consists of $p^s_{12} = d_1 + \mu_2(c_1)$, $p^s_{22} = \mu_d(d_1) + \mu_2(c_1)$, *and*

$$p^s_{11} = \begin{cases} p^*_{11} & \text{if } p^*_{11} \geq \mu_d + \mu_c \text{ and} \\ & \quad \pi^s_b(p^*_{11}) \geq \pi^s_b(\mu_d + \mu_c) \\ \mu_d + \mu_1 & \text{otherwise,} \end{cases}$$

where $\pi^s_b(p_{11})$ *is given by,*

$$\pi^s_b(p_{11}) = 2s - p_{11} - \mu_2 - \Pr(c_1 + d_1 > p_{11}) \times \mathbb{E}(\mu_d(x) + k \mid c_1 + d_1 > p_{11}) - \Pr(c_1 + d_1 \leq p_{11}) \times \mathbb{E}(\min(d_1, \mu_d(x) + k) \mid c_1 + d_1 \leq p_{11}). \quad (3)$$

The buyer switches suppliers in period 2 *if* $k + \mu_d(d_1) < d_1$.

The key to understanding the form of the optimal contract in Lemma 1 lies in understanding the shape of the expected profit function in Figure 4. The slope of $\pi^s_b(p_{11})$ asymptotically approaches -1 as p_{11} goes to $\pm\infty$, and π^s_b has a convex–concave shape. Thus, $\pi^s_b(p_{11})$ either has a local maximum (which is p^*_{11} from the lemma) or is decreasing everywhere. Ensuring that supplier 1's participation constraint holds limits the buyer to a feasible region consisting of $p_{11} \geq \mu_d + \mu_1$. Thus, the optimal first-period price is either $\mu_d + \mu_1$ or p^*_{11}.

The intuition behind this result is that, if switching costs are small, the buyer's profit is likely to be decreasing in p_{11}. There is little consequence to failure; hence the buyer offers the lowest possible price. However, if switching costs are high, $\pi^s_b(p_{11})$ has a shape like that depicted in Figure 4. If p_{11} is very small, increasing it slightly does nothing to lower the chance of default and only costs the buyer more. Likewise, if p_{11} is very large, then a slight increase does little to affect the probability of default. However, if p_{11} is intermediate in value then a small change may result in a large decrease in the probability of failure, outweighing the excess cost to the buyer. Thus, it may be optimal to offer a price that is higher than the supplier's minimum acceptable price to lower the probability of default.

Figure 4 An Example of the Buyer's Expected Profit as a Function of p_{11} in the Single-Period Contract with the Risk of Failure

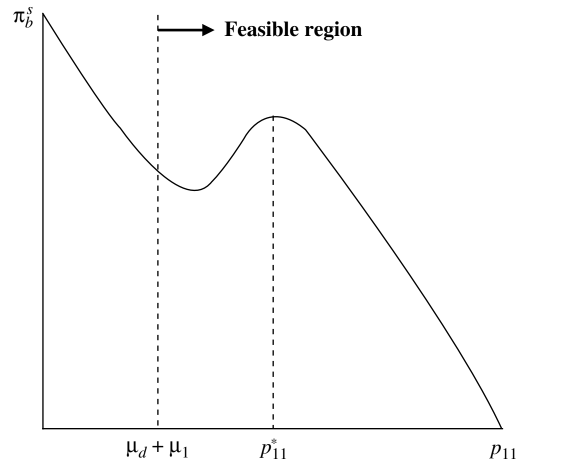

It is interesting that the optimal contract derived in Lemma 1 is identical to the optimal contract in Theorem 1, except for the first-period price p^s_{11}, because failure is irrelevant (from the buyer's point of view) in the second period. Furthermore, with failure, the first-period price is greater than the first-period price without failure, because the minimum possible price from Lemma 1 is equal to the optimal price in Theorem 1. Thus, comparing short-term contracts, we see that the buyer pays more when suppliers face an endogenous default risk than when suppliers have no risk of default, to reduce the probability of failure and hence the chance of incurring the switching cost k. In addition, by applying the Envelope Theorem, it can be shown that the buyer's optimal profit is decreasing in k. Intuitively, the more expensive it is to switch suppliers in the middle of the product's sale horizon, the lower the buyer's expected profit.

5.2. Long-Term Contract

In a long-term static contract, the buyer offers a fixed set of prices $\{p_{11}, p_{12}\}$ to the first supplier. If the supplier accepts the contract, the buyer switches suppliers only in the event of supplier 1's bankruptcy. Thus, in contrast to the short-term contract, there are two rather than three regions of interest. In Region I, supplier 1 survives the first period, and the buyer continues to do business with that supplier at the agreed upon price of p_{12} in the second period. In Region III, supplier 1 fails, leaving the buyer with only one alternative: to switch to supplier 2. See Figure 5 for a graphical representation of these regions. The slope of the line is fixed and cannot be controlled by the buyer, who is thus incapable of replicating the optimal switching (found in Figure 2) if supplier 1 turns out to be high-cost.

In Lemma 2, it will be useful to define the following function: Let $p_{12}(p_{11})$ be the second-period price when the first-period price is p_{11} and supplier 1's participation constraint binds.

Lemma 2. *Let p_{11}^l be the solution to*

$$\int_{-\infty}^{\infty}(\mu_d(x)+k-x)f(p_{11}^l-x)g(x)\,dx=0. \tag{4}$$

Figure 5 Optimal Second-Period Action of the Buyer as a Function of the Realized Values of c_1 and d_1 Under the Long-Term Contract

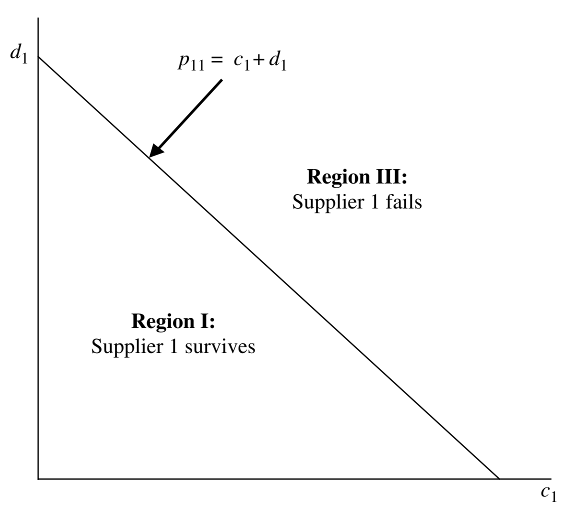

Then, the optimal long-term contract under the threat of default is p_{11}^l, $p_{12}^l = p_{12}(p_{11}^l)$, and $p_{22}^l = \mu_d + \mu_2(c_1)$.

Note that in Lemma 2 we have not restricted p_{12}^l to be non-negative. Numerically, it is rare to observe $p_{12}^l < 0$ but not impossible. For this to occur, the switching cost must be very large (e.g., an order of magnitude greater than the average total production cost). The economic interpretation of a negative second-period price is that, if the expected cost incurred due to default is extremely large (i.e., if the chance of default is high or k is large), it is optimal for the buyer to heavily subsidize the supplier. In return, if costs turn out to be low, the supplier reimburses the buyer in the second period for insuring the firm against default.

5.3. Contract Choice

To determine which contract the buyer prefers, we must compare expected profits under the optimal contract in each case.

Theorem 2. *In the presence of failure risk,* (i) π_b^s, $\pi_b^l \le \bar{\pi}_b$ *and* (ii) *there exists some k^* such that, for all $k > k^*$, $\pi_b^s \le \pi_b^l$.*

In other words, the long-term contract is preferred to the short-term contract if switching costs are high, but neither contract achieves the system optimal profit $\bar{\pi}_b$. The first part of Theorem 2 is intuitive, because we expect the system to perform no better under the threat of default than a system without failure. The second part of the theorem demonstrates that long-term contracts are preferred when switching costs are high. Essentially, long-term contracts allow the buyer to shift more of the total compensation to the first period, thus lowering the chance of supplier failure. This comes at the expense of losing the option to voluntarily switch suppliers, and hence the buyer prefers the long-term contract only if the savings due to the decreased chance of default outweigh the lost option to switch suppliers. Numerically, we observe that the threshold k^* is typically very small in relation to the average production costs in the system (e.g., an order of magnitude or more), and in many cases the long-term contract is preferred for all non-negative values of k. In a numerical study consisting of 243 sets of parameters (see §5 of the appendix, which is available online, for details), we found that, on average, k^* was 27% of the total mean per unit production cost. Thus, for many reasonable parameters (i.e., moderately significant switching costs satisfying $k \gtrsim 0.3(\mu_c + \mu_1)$), the buyer prefers the long-term contract.

It is also interesting to note how the value of k^* changes as a function of the variability in the idiosyncratic cost parameter. Figure 6 provides an example of the typical behavior observed in numerical experiments using normally distributed costs. First, we note that, for any given value of the coefficient of variation, the value of k^* is decreasing in ρ_d. This is due to the fact that the value of the option to voluntarily switch suppliers is also decreasing in ρ_d. Hence the relative value of the short-term contract compared to the long-term contract is decreasing as costs become more correlated. Thus, the buyer is more likely to prefer the short-term contract for smaller switching costs. In addition, for fixed ρ_d, the threshold k^* has a quasi-concave shape (although we note that, depending on the problem parameters, the functions do not necessarily have a decreasing portion). This is due to the fact that there are two competing forces at play affecting the value of the contracts. As the variability of d_i increases, the value of the option to switch suppliers in the short-term contract increases. The buyer is able to exploit the low realizations of d_i more efficiently in the short-term contract; hence, k^* is increasing. On the other hand, as the variability of d_i increases, the chance of default also increases, which increases the relative attractiveness of the long-term contract. If the variability is very high, failure is frequent and costly; hence, the default effect dominates the option effect and k^* is decreasing. The opposite holds if the variability is very low because the chance of failure is small.

Figure 6 An Example of the General Behavior of the Threshold Switching Cost as a Function of Variability in the Idiosyncratic Production Cost Observed in Numerical Experiments

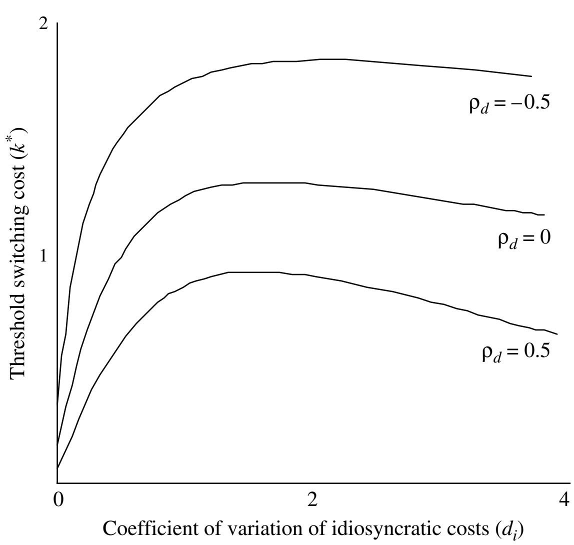

Note. Both idiosyncratic and common costs are normally distributed with mean 3.

6. Dynamic Contracts

Now that we have shown that the previously described contracts perform worse than the centralized system, we move on to a class of contracts that coordinate the channel. This class of contracts is dynamic or state-dependent, as opposed to the previous static contracts. The sequence of events is the same. The difference between the two is that the prices in a static contract are fixed, whereas prices in a dynamic contract may be tied to some index to help insulate the supplier against failure by shifting some of the risk to the buyer. For example, suppose that the uncertainty in the common cost c_t is primarily due to fluctuations in the global petroleum market. In forming a dynamic contract, the buyer might tie the contract price to the commodity price of oil, compensating the supplier for part or all of the variation in the common cost. These types of contracts are frequently observed in the Japanese industry, contrasting with traditionally static contracts in the U.S. industry. For example, McMillan (1990) describes how Japanese manufacturers typically do not specify prices in initial contracts, but rather update prices every six months based on a review of the supplier's production costs, considering separately such issues as labor, raw materials, design changes, and energy costs. Buyers typically allow unavoidable cost increases, such as raw materials, to be reflected in the contract price but are less likely to allow increases due to controllable costs, such as labor.

For the purpose of our analysis we assume that the dynamic contract compensates the supplier perfectly for c_t. Effectively, this assumption removes c_t from the contracting problem by having the buyer bear this cost in its entirety. The buyer must then decide how to compensate the supplier for idiosyncratic costs, d_i. Because this contract will be shown to coordinate the channel, there is no loss of generality in restricting attention to this particular form of cost compensation as opposed to some other form (e.g., dynamically compensating for idiosyncratic costs or some combination of the two components).

We assume that there is no additional cost to implementing a dynamic contract. Yet dynamic contracts may be difficult or costly to administer, which may reduce their relative attractiveness. Administrative costs are not explicitly modeled here but might include, for example, tracking and verification of the pricing index that determines the supplier's compensation.

6.1. Dynamic Contracts Without Failure

The following lemma details the optimal contracts for the no-failure case.

Lemma 3. (i) *The optimal short-term dynamic contract is* $p^s_{11} = \mu_d + c_1$, $p^s_{12} = d_1 + c_2$, *and* $p^s_{22} = \mu_d(d_1) + c_2$. *The buyer switches suppliers in the second period if* $k + \mu_d(d_1) < d_1$.

(ii) *The optimal long-term dynamic contract is any pair* $\{p^l_{11}, p^l_{12}\}$ *such that* $p^l_{11} + p^l_{12} = 2\mu_d + \mu_1 + \mu_2$.

(iii) *The expected profit in each dynamic contract is equal to the expected profit in their static counterparts in Theorem* 1.

Lemma 3 provides an interesting result: In the absence of failure risk, in terms of expected profit, the dynamic contracts are equivalent to static contracts of the same length. In other words, to a risk-neutral buyer, choosing a dynamic contract offers no advantage.

6.2. Dynamic Contracts with Failure

Dynamic contracts transfer the risk of common cost uncertainty from the supplier to the buyer, thus lowering the probability of failure due to high common costs. Because we assume that the buyer is a large, risk-neutral firm, transferring this risk increases the relative attractiveness of the dynamic contracts to the buyer when supplier failure is a possibility. The following lemma demonstrates this.

Lemma 4. (i) *The optimal long-term dynamic contract is* $p^l_{11} = \alpha + c_1$, $p^l_{12} = p_{12}(p^l_{11})$, *and* $p^l_{22} = \mu_d(d_1) + c_2$, *where* $p_{12}(p^l_{11})$ *is the dynamic second-period price for which the supplier's participation constraint binds*.

(ii) *The optimal short-term dynamic contract is given by* $p^s_{11} = c_1 + \max(x^*, \mu_d)$, $p^s_{12} = d_1 + c_2$, *and* $p^s_{22} = \mu_d(d_1) + c_2$, *where* x^* *is the solution to*

$$-1 + g(x^*)(-x^* + \mu_d(x^*) + k) = 0.$$

(iii) *The long-term dynamic contract is preferred to the short-term dynamic contract and yields expected profit equal to the system optimal expected profit without failure risk*.

It is interesting that in Lemma 4 we have precisely the opposite result from Theorem 1: Long-term dynamic contracts are always preferred to short-term dynamic contracts. This is because long-term contracts allow the buyer to switch suppliers in the optimal manner; by setting $p_{11} = \alpha + c_1$, supplier 1 fails (and hence the buyer switches suppliers) if $\mu_d(d_1) + k < d_1$, exactly the same action that the centralized system would take. The second-period price then serves as a compensation mechanism to ensure that supplier 1 has a binding participation constraint. In the short-term contract, however, second-period prices must be subgame perfect. Thus, the buyer cannot promise a price that justifies setting $p_{11} = \alpha + c_1$ as the first-period price, and hence cannot switch suppliers in the system optimal manner. Consequently, the first-period price is strictly less than $\alpha + c_1$, and total profits are lower than in the long-term contract.

The fact that long-term dynamic contracts in the presence of default risk achieve the system optimal solution *without failure* is intriguing. It means that buyers can simultaneously lock in suppliers and form lasting relationships built on trust and dynamic cost compensation, yet still achieve the proper filtering of costly suppliers to maximize their own profits. Indeed, if we were to plot the actions resulting from the optimal dynamic long-term contract it would look exactly like Figure 2, with the only difference being that the regions in the dynamic contract case are determined by failure rather than by the buyer's choice to switch suppliers. Furthermore, because a first-period price of $\alpha + c_1$ is never optimal in the short-term dynamic contract, this contract is incapable of replicating the switching in Figure 2 and hence cannot coordinate the channel.

The value of a dynamic contract lies in its ability to remove the stochastic element that affects both suppliers from the factors leading to default. With static contracts, a potentially efficient supplier (i.e., a supplier with lower idiosyncratic costs than the expected costs of the alternative supplier) may fail due to high common costs, which reduce overall supply chain efficiency. With dynamic contracts, on the other hand, suppliers fail only if their idiosyncratic costs are high, precisely the situation in which the buyer would voluntarily switch suppliers. Hence, dynamic contracts eliminate a harmful source of stochasticity (common costs) while retaining a potentially useful source of stochasticity (idiosyncratic costs), from the supply chain's point of view.

7. Extensions

In this section we analyze four independent extensions to the core model that allow us to comment further on the scope of our results. In the first subsection we consider the effects of demand uncertainty. The second subsection addresses contingent transfer payments—that is, transfer payments made from the buyer to a supplier (possibly dependent on cost realizations) intended to subsidize the supplier and prevent bankruptcy. The third subsection addresses renegotiation. The final subsection discusses the special case of normally distributed costs, using the increased specificity of the model to answer several interesting questions concerning the performance of the various contracts as a function of cost correlation and uncertainty.

7.1. Demand Uncertainty

We have assumed throughout the paper that the buyer's demand is deterministic and equal to one in each period. In practice, suppliers may face demand uncertainty in addition to cost uncertainty. However, because we have explicitly incorporated cost uncertainty into the model, this assumption is equivalent to a make-to-order (MTO) system with uncertain demand. For example, consider a supplier that only produces units that are purchased by the buyer. The total size of the buyer's order in period t is a random variable D_t that is independent of both idiosyncratic costs, d_i, and common costs, c_t. The supplier may face exogenous capacity constraints, in which case D_t is truncated at the capacity level of the supplier with mass added to the endpoint. The profit to supplier i in period t is then $D_t(p_t - d_i - c_t)$.

Assume that the supplier must make a minimum profit of K_t in period t to survive. K_t may represent the supplier's annual (fixed) operating costs, interest payments on outstanding loans, capital outlay for the new production process, legacy pension expenses, etc. Note that K_t is allowed to be negative; i.e., the supplier may be allowed to lose some money and still survive and may evolve across time. However, because suppliers are ex ante identical, K_t is the same for both suppliers in each period. Thus, supplier i survives in period t if $D_t(p_t - d_i - c_t) \geq K_t$. Because all the random variables are independent of one another, we may write this as $p_t - d_i - c_t \geq K_t/D_t$. Redefining the common cost to incorporate the demand term, $c'_t = c_t + K_t/D_t$, we see that supplier i survives if and only if profit is nonnegative, i.e.,

$$p_t - d_i - c'_t \geq 0. \tag{5}$$

In addition, we may consider the case in which the supplier only accepts a contract that yields expected profits of at least K_t. In that case, the same logic yields the result that the supplier accepts any contract in which (5) holds in expectation. Both of these conditions—supplier survival and contract acceptance—are identical to the case with deterministic demand in each period, so long as the common cost terms are properly defined. Consequently, our assumptions of deterministic demand and no fixed operating expenses are made without loss of generality in an MTO production system with exogenous capacity constraints.

This equivalence also yields insight into why the common cost term is not diversifiable and hence cannot be hedged. Because the common cost term may be thought of as capturing demand risk as well as common cost risk, it is unlikely that this risk could be mitigated. Furthermore, because the cost term also depends on materials prices that cannot be procured from any industrial exchange (e.g., the outputs of upstream suppliers), our implicit assumption that cost uncertainty cannot be hedged is justified.

It is important to note that our extension to the case of demand uncertainty is valid only if the supplier is uncapacitated or faces an exogenous capacity constraint. This scenario reflects our example (the auto industry) in which the capacity is sometimes dictated by the buyer, suppliers may not have sufficient capital for creating excess capacity, capacity constraints on components are typically not tight (i.e., it is more likely that assembly capacity is binding), and, if capacity is likely to be tight, components are multisourced. By focusing on a single-sourced component, we have implicitly assumed that capacity constraints are not an issue for the supplier. A model in which capacity is tight may be more suitable to an analysis of multisourcing (see Tomlin 2006). In any event, endogenous capacity decisions are likely to provide further reasons to favor long-term contracts, because long-term relationships are known to stimulate capacity investment (Taylor and Plambeck 2007). Hence it is unlikely that endogenizing the capacity decision would significantly alter our results.

7.2. Contingent Transfer Payments and Loans

The contract types we have discussed thus far are fairly simple, consisting of either fixed prices or prices that are a function of one of the stochastic elements (i.e., the common cost). When a supplier enters bankruptcy, we provide no recourse to the buyer to help alleviate the supplier's financial distress. This may be an acceptable assumption if the buyer is unable or unwilling to subsidize the supplier in the event of bankruptcy (e.g., if capital is expensive to the buyer or if the buyer is also in financial distress). However, in some situations the buyer may prefer to make a transfer payment that allows the supplier to avoid bankruptcy and continue operations. For example, the Big Three Detroit automakers provided $100 million in direct operating subsidies to support Collins & Aikman in bankruptcy in 2006 (Barkholz and Sherefkin 2007). In this section we discuss this scenario.

We consider the following modification of the core model: If the first supplier enters bankruptcy, then at the start of the second period, upon observing the realized value of all costs, the buyer has the option to either switch suppliers or make a transfer payment to the first supplier, thus raising the capital level to zero and preventing bankruptcy. In making this decision, the buyer takes into account the size of the necessary transfer payment as well as the costs of contracting with each supplier and any switching costs. Because resulting transfer payment depends on the realized value of *both* first-period cost terms, it is termed a contingent transfer payment. In addition to allowing direct operating subsidies that are not reimbursed, we consider loans from the buyer to the supplier that may be partially or fully repaid (with interest). We refer to either scenario (no repayment or repayment) as a transfer payment.

For the details of our model of transfer payments, see §2 of the technical appendix. Note that we make the following critical assumption about the timing of the payment: Any transfer payment occurs at the start of the second period and such payments are not considered in the short-term (i.e., one period) contract participation constraint of the supplier. The implication of this assumption is that the supplier is unwilling to accept a lower first-period price if the buyer has the option of offering a transfer payment when bankruptcy occurs. The short-term participation constraint represents the outside option of the supplier for the immediate future. Given the financial distress of the firm, the primary concern (particularly when engaging in a short-term contract without the promise of future business) is likely short-term financial performance. Hence the supplier is likely to be unwilling to sacrifice short-term profit for a potential future transfer payment, especially when this may greatly increase the chance of bankruptcy. The consequences of this assumption are discussed in more detail below.

It is clear that the buyer does at least as well when allowed to make a transfer payment as in the simpler contracts discussed previously, because the transfer payment is entirely optional. Thus, contracts with transfer payments are preferred to contracts without such payments. The questions we then seek to address are the following. First, does the result of Theorem 2 hold with transfer payments? In other words, are long-term contracts preferred when switching costs are high, even if transfer payments are available under both short- and long-term contracts? Second, does either contract type coordinate the system (i.e., achieve the first best solution that the long-term dynamic contract without transfer payments achieves)?

The following theorem answers both of the questions, essentially demonstrating that all of the results of Theorem 2 hold even when the buyer is given the option of subsidizing distressed suppliers.

Theorem 3. *If the buyer is allowed to make contingent transfer payments or loans, then* (i) *neither the short-term nor the long-term contract achieves the first best profit* ($\bar{\pi}_b$) *and* (ii) *there exists some* k^* *such that, for all* $k > k^*$, *the long-term contract is preferred to the short-term contract.*

The intuition behind this result is that transfer payments and loans allow the buyer to avoid switching suppliers in the event of failure, provided the incumbent supplier is efficient enough to warrant a subsidy. However, this benefit of transfer payments applies to both short- and long-term contracts in different ways. Transfer payments are valuable with short-term contracts because these contracts typically involve lower first-period prices and higher failure rates than long-term contracts. Hence, the recourse provided in subsidizing a bankrupt supplier is greater with short-term contracts because failure happens more often. On the other hand, transfer payments introduce an option to stay with an efficient (albeit bankrupt) supplier in long-term contracts. Hence, there is an increase in value due to the second-period option effect. Numerically, neither benefit dominates, and the effect on k^* (compared to the case of no transfer payments) is ambiguous: In a numerical study consisting of 243 sets of parameters (see §5 of the appendix, available online, for details), we found that, on average, k^* was 18% of the total mean per unit production cost, compared to 27% without transfer payments. In 42 of 243 case k^* increased due to the presence of transfer payments (i.e., the transfer payment provided more value to the short-term contract than to the long-term contract). In the remaining 201 cases k^* was lower with a transfer payment.

Recall that we assumed that any transfer payment is made at the start of the second period (or, more important, that the supplier does not consider the transfer payment in his short-term contract participation constraint). This is a key assumption: If the supplier takes the transfer payment into account, for large k the short-term contract essentially transforms into a long-term contract in the following sense. If switching costs are high, the buyer pays the supplier nothing for first-period production. Consequently, the supplier always fails. The buyer never switches suppliers, however, preferring instead to make a transfer payment to the bankrupt supplier and avoid high switching costs. The short-term contract has effectively become a long-term contract in which the buyer promises to work with supplier 1 in the second period and pay all production costs *ex post* after the supplier enters bankruptcy. As a result the expected profit is the same with a long-term and short-term contract as k becomes large; hence, there is not a strict preference between the two. (For low switching costs, as in all other cases that we have considered, the short-term contract is preferred to the long-term contract.)

This equivalence is eliminated and a strict preference for long-term contracts is restored if any of several complications arise, including the following: if the supplier is unwilling to accept certain bankruptcy (because of a contract price of zero) in the first period, i.e., there is a minimum acceptable contract price; if there is some fixed cost to making a transfer payment (e.g., the buyer must pay the supplier's bankruptcy or default penalty); or if the supplier discounts future revenues (which implies that the supplier will not accept a contract price of zero even if a transfer payment is made in the future). Thus, although the timing of the transfer payment is important to the result of Theorem 3, this assumption may be relaxed while preserving the result if one of a variety of alternative conditions holds.

Although we have assumed that the suppliers in the current analysis are incapable of securing external funds in the event of bankruptcy between periods one and two, it is interesting to consider this case. If the supplier is capable of borrowing enough funds in any scenario to avert bankruptcy, then clearly failure has no effect on the buyer. The supplier always avoids bankruptcy, and hence the model is equivalent to the model without default. If the supplier is capable of borrowing limited funds, however, and the chance of default remains, then the core results of the model are preserved. It is still true that long-term contracts hold value in allowing the buyer to rearrange the cash flows and helping the supplier avoid default.

7.3. Renegotiation

We explicitly excluded the possibility of supplier renegotiation (i.e., supplier hold-up) in long-term contracts due to the strong bargaining power of the buyer. However, a *buyer*-initiated renegotiation is possible, as in the Ford example discussed in the introduction. It can be shown (see §3 of the technical appendix, available online) that, if the buyer is allowed to renegotiate a long-term contract in the second period, all of the results are preserved. This result critically depends on the fact that, with or without renegotiation, the supplier's participation constraint is binding in the optimal contract. Thus, the buyer extracts all surplus from the supplier, and, if renegotiation occurs (and the supplier anticipates the renegotiation) the buyer must compensate the supplier via a higher first-period price to satisfy the supplier's participation constraint. In other words, any additional funds extracted via renegotiation must be compensated for via the contract price. The net result is that the buyer's expected profit is the same regardless of whether renegotiation occurs; hence, the preference between contracts remains identical to the cases already discussed in §§3–6.

Interestingly, if the supplier does not anticipate renegotiation in a long-term contract (i.e., if renegotiation is not taken into account in the supplier's participation constraint), then the buyer enjoys strictly greater profits in a long-term contract than in a model without renegotiation because the buyer need not compensate the supplier with a higher contract price. As a result, long-term contracts have even greater value than previously discussed, and are thus preferred for large switching costs. Hence, the results of the paper hold even when the buyer is allowed to unilaterally renegotiate long-term contracts, regardless of whether the supplier anticipates renegotiation.

7.4. Normally Distributed Costs

By assuming that costs are normally distributed, we derive further insights into the behavior of the various contracts. In what follows we consider three contract types in the presence of default risk: the long-term static and dynamic contracts and the short-term static contract. Because the long-term dynamic contract dominates the short-term dynamic contract, the latter is omitted. Let c_1, c_2 be identically distributed (possibly correlated) $N(\mu_c, \sigma_c)$ random variables, and let d_1, d_2 be bivariate normal with identical mean and variance μ_d and σ_d^2 and correlation ρ_d. From the properties of the bivariate normal distribution, the expected value of d_2 conditional on $d_1 = x$ is $\mu_d(x) = (1-\rho_d)\mu_d + \rho_d x$. Hence, it is optimal to switch suppliers in the second period if $d_1 > \mu_d + k/(1-\rho_d)$. Because α is increasing in ρ_d, the buyer seems less likely to switch suppliers if costs are highly correlated. The following theorem further describes the behavior of the contracts as a function of ρ_d and σ_d.

Theorem 4. (i) *The optimal expected profit under all contract types is decreasing in ρ_d.*

(ii) *The difference between the system optimal (long-term dynamic) profit and the profit under the long-term static contract is decreasing in ρ_d. In the limit as $\rho_d \to 1$, profits are equal.*

(iii) *The centralized system optimal expected profit is increasing in σ_d.*

Intuitively, from part (i), system profit is lower if the idiosyncratic costs of the two firms are highly correlated. There is little value in the option to switch suppliers, so the overall expected profit of the buyer is higher when suppliers have negatively correlated idiosyncratic costs.

Part (ii) demonstrates how the relative advantage of the dynamic long-term contract varies as a function of ρ_d. If costs are strongly negatively correlated, then the dynamic long-term contract performs quite well. In this case, the dynamic contract is effective at switching suppliers in the optimal manner, whereas the static contract is less efficient. If costs are strongly positively correlated, however, the value of switching suppliers is lower, hence the performance gap between the two contract types is much smaller, although dynamic contracts offer value for any $\rho_d < 1$. Consequently, long-term dynamic contracts are most valuable in situations where suppliers have negatively correlated costs.

Part (iii) describes the behavior of the system optimal expected profit as a function of variability in the suppliers' private cost. Intuitively, the more variable the suppliers' costs, the more likely a low cost realization. The buyer is shielded from high cost realizations by the option to switch suppliers. Thus, for fixed ρ_d and μ_d, increased variability in the suppliers idiosyncratic costs allows the buyer to exploit the option to switch.

8. Discussion

In this paper we have presented a model of buyer-supplier contracting primarily characterized by three features: uncertain production costs, extended sales horizons, and the strong bargaining power of the buyer. Within this context we have shown that the threat of supplier failure can increase the buyer's preference for long-term contracts. Furthermore, dynamic contracts that compensate suppliers for common costs (materials, etc.) achieve the system optimal profit. This feature helps to explain the adoption of these contracts in the Japanese auto industry (McMillan 1990).

We did not model risk-averse firms because discussing supply chain coordination in such a setting adds another level of complexity that is outside the scope of our work (Gan et al. 2004). However, the true value of dynamic contracts may depend on the risk-neutrality (or lack thereof) of the buyer. In practice, auto manufacturers do not always reimburse suppliers for the full amount of common costs. One potential explanation for this is that the buyer is not risk-neutral and hence does not seek to bear all of the risk associated with the variability in raw materials costs. Another possible explanation may be that the variability associated with c_t is not completely outside of the supplier's control; hence, the buyer needs to leave some of the risk with the supplier to induce the proper actions (e.g., negotiating low prices from second-tier suppliers, etc.). Finally (and perhaps most important), dynamic contracts are difficult to administer and are significantly less formal than static contracts, and trust between partners is key to their implementation. There has historically been a severe lack of trust between U.S. auto manufacturers and suppliers, perhaps helping to explain why both parties are hesitant to engage in dynamic contracting agreements.

Throughout the analysis we have ignored the effects of learning curves, seen for example in Spence (1981). In much of the multiperiod contracting literature, learning curves are modeled as cooperative improvements in design and production processes that reduce costs in long-term relationships (Cohen and Agrawal 1999), whereas short-term relationships offer fewer (or no) opportunities for cost reduction. If such a learning curve were present in our model, it would serve to increase the profitability of the long-term contract, thus providing further incentive for a buyer to choose this contract type. Qualitatively, our main results should be strengthened by this complication. In addition, considering asymmetric suppliers yields a similar result: If one supplier's costs dominate the other supplier, the buyer will seek to contract in the first period with the more efficient supplier. The long-term contract then serves as a tool to both mitigate switching costs and prevent switching to a less efficient supplier; thus, the value of the long-term contract is increased.

An interesting counterpart to the case of contingent transfer payments are contracts that explicitly allow the buyer to end the relationship should the production costs of the supplier exceed a certain threshold (i.e., the buyer is provided a means of breaking the relationship if a supplier is very inefficient). Such a contract, which would provide the buyer the option of switching suppliers if the first supplier survives (rather than fails) increases the profitability of long-term relationships and hence increases the relative attractiveness of a long-term contract compared to a short-term contract, leaving our main results intact. Such contracts are unable to coordinate the system, however, because of the forced switching of suppliers that occurs if the incumbent supplier fails. Thus, dynamic contracts still perform better by achieving coordination.

There are two complications that may increase the relative value of short-term contracts: discounting of the buyer's second-period profit, and non-zero bankruptcy costs for the suppliers. The first complication lessens the impact of supplier default and the cost of voluntary switching: Because both of these costs are incurred in the second period, discounting decreases their relative contribution to the total expected profit. The second complication effectively increases the switching costs due to default while leaving unchanged the switching costs due to voluntary changes in the second-period supplier. Intuitively, if the suppliers incur some bankruptcy penalty, the buyer must compensate the supplier more to satisfy the participation constraint. Thus, if the participation constraint is binding, the buyer's profit function includes an extra term penalizing for the bankruptcy costs if default occurs, which is essentially the same as increasing the switching cost k. The extra term is not present, however, when the supplier survives but the buyer voluntarily switches suppliers; hence, the relative value of the short-term contract is increased because voluntary switching is less costly. Both of these complications essentially increase the threshold k^* from Theorem 2, but other results remain qualitatively unchanged.

Our model is one of partial equilibrium. In reality, as distressed suppliers declare bankruptcy and exit the market, new suppliers may enter, perhaps in better financial standing. This effect is important in the automotive industry but occurs over a long time (e.g., it may take years for a newly created supplier to build the necessary technology to produce at the scale and quality level of a large, existing supplier such as Delphi). Our model analyzes the medium-term issue of dealing with a distressed supplier base. In this time frame it is often impractical for a buyer to work with an emerging supplier: Particularly in the automotive industry, asset-specific investments in suppliers are quite large and hence switching costs are prohibitive.

With no information asymmetry for suppliers to use as leverage, we have created a situation wherein buyers extract all of the surplus in the supply chain. This assumption has provided the strongest incentive for the buyer to engage in a short-term contract to exploit the option to voluntarily switch suppliers. Still we find that long-term commitment is very often the more profitable choice; unlike the results in Li and Debo (2007), the attractiveness of the long-term contract in our model is not driven by increased price competition in the supplier's bidding strategies, but rather by the decreased chance of failure resulting from a long-term commitment.

The main conclusion of our analysis is that long-term contracts may be even more valuable to buyer-supplier relations—particularly within the American automotive industry—than previously thought. The literature abounds with reasons for firms to prefer long-term contracts: as a manner of developing trust and cooperation between partners (Dyer 1996), as a tool to increase price competition in auction scenarios (Li and Debo 2007), and as relational tools to enforce actions that are otherwise unsupportable or uncontractable in short-term situations (Taylor and Plambeck 2007). We add to those reasons by demonstrating the value of long-term relationships when suppliers face the threat of failure.

Electronic Companion

An electronic companion to this paper is available on the *Manufacturing & Service Operations Management* website (http://msom.pubs.informs.org/ecompanion.html).

Acknowledgments

The authors thank Vlad Babich, Chris Tang, the associate editor, and three anonymous referees for many helpful comments that greatly improved the paper.

References

Andersson, B. 2006. GM's global supply footprint. Presentation, The Supplier Industry in Transition Conference, April 19, Federal Reserve Bank of Chicago, Detroit, MI.

Archibald, T. W., L. C. Thomas, J. M. Betts, R. B. Johnston. 2002. Should start-up companies be cautious? Inventory policies which maximise survival probabilities. *Management Sci.* **48**(9) 1161–1174.

Atkins, D., H. Krishnan, X. Zhao. 2005. Relational price-only contracts in supply chains. Working paper, University of British Columbia, Vancouver, BC, Canada.

Babich, V. 2006. Dealing with supplier bankruptcy: Costs and benefits of financial subsidies. Working paper, University of Michigan, Ann Arbor, MI.

Babich, V., A. N. Burnetas, P. H. Ritchken. 2007. Competition and diversification effects in supply chains with supplier default risk. *Manufacturing Service Oper. Management* **9**(2) 123–146.

Barkholz, D., R. Sherefkin. 2006. Lear left reeling from Dodge Ram snub. *Automotive News* **80**(6187) 1–48.

Barkholz, D., R. Sherefkin. 2007. C&A debacle will cost automakers $665 million. *Automotive News* **81**(6248) 1–47.

Cachon, G. 2003. Supply chain coordination with contracts. *Handbooks in Operations Research and Management Science: Supply Chain Management*, Chap. 6. North Holland, Amsterdam, The Netherlands.

Choi, T. Y., J. L. Hartley. 1996. An exploration of supplier selection practices across the supply chain. *J. Oper. Management* **14** 333–343.

Cohen, M. A., N. Agrawal. 1999. An analytical comparison of long and short term contracts. *IIE Trans.* **31**(8) 783–796.

Dyer, J. 1996. How Chrysler created an American keiretsu. *Harvard Bus. Rev.* **74**(4) 42–56.

Dyer, J., D. S. Cho, W. Chu. 1998. Strategic supplier segmentation: The next "best practice" in supply chain management. *California Management Rev.* **40**(2) 57–77.

Gan, X., S. P. Sethi, H. Yan. 2004. Coordination of supply chains with risk-averse agents. *Production Oper. Management* **13**(2) 135–149.

Kleindorfer, P. R., D. J. Wu. 2003. Integrating long- and short-term contracting via business-to-business exchanges for capital-intensive industries. *Management Sci.* **49**(11) 1597–1615.

Levy, E., S. Ferazani. 2006. Standard & Poor's industry survey: Autos and auto parts. *Standard Poor's* **174**(51) 1–39.

Lewin, T. 2006. U.S. suppliers count on Europe for rescue. *Automotive News* **80**(6199) 22D.

Li, C., L. Debo. 2007. Strategic dynamic sourcing from competing suppliers: The value of commitment. Working Paper E2005E20, Tepper School of Business, Carnegie Mellon University, Pittsburgh, PA.

McMillan, J. 1990. Managing suppliers: Incentive systems in Japanese and U.S. industry. *California Management Rev.* **32** 38–55.

Nussel, P., D. Barkholz. 2006. Delphi demands price hikes. *Automotive News* **80**(6196) 1–38.

Plambeck, E., T. Taylor. 2007. Implications of renegotiation for optimal contract flexibility and investment. *Management Sci.* **53**(12) 1872–1886.

Rudambi, R., S. Helper. 1998. The "close but adversarial" model of supplier relations in the U.S. auto industry. *Strategic Management J.* **19** 775–792.

Sako, M., S. Helper. 1998. Determinants of trust in supplier relations: Evidence from the automative industry in Japan and the United States. *J. Econom. Behavior Organ.* **34** 387–417.

Sherefkin, R. 2006. Ford seeks steel futures market. *Automotive News* **81**(6228) 1–48.

Spence, A. M. 1981. The learning curve and competition. *Bell J. Econom.* **12**(1) 49–70.

Srinivasan, R., T. H. Brush. 2006. Supplier performance in vertical alliances: The effects of self-enforcing agreements and enforceable contracts. *Organ. Sci.* **17**(4) 436–452.

Stallkamp, T. T. 2005. *Score! A Better Way to Do Business*. Wharton School Publishing, Upper Saddle River, NJ.

Swinney, R., G. Cachon, S. Netessine. 2005. Capacity investment by competitive start-ups. Working paper, University of Pennsylvania, Philadelphia, PA.

Tang, C. S. 1999. Supplier relationship map. *Internat. J. Logist.: Res. Appl.* **2**(1) 39–56.

Taylor, T., E. Plambeck. 2007. Supply chain relationships and contracts: The impact of repeated interaction on capacity investment and procurement. *Management Sci.* **53**(10) 1577–1593.

Tomlin, B. 2006. On the value of mitigation and contingency strategies for managing supply chain disruption risks. *Management Sci.* **52**(5) 639–657.

Tunca, T. I., S. A. Zenios. 2006. Supply auctions and relational contracts for procurement. *Manufacturing Service Oper. Management* **8**(1) 43–67.

Wernle, B. 2006a. Suppliers in "shambles," so why's Ross buying? *Automotive News* **80**(6199) 28.

Wernle, B. 2006b. Toyota: Supplier-automaker ties need shoring up. *Automotive News* **80**(6199) 4.

Wynn, G. 2006. General Motors woes seen causing supply chain cull. *Reuters Transportation* (January 9).